arXiv:1608.04961v2 [stat.ML] 7 Oct 2016

Clustering Mixed Datasets Using Homogeneity Analysis with Applications to Big Data

Rajiv Sambasivan *

October 10, 2016

Abstract

Clustering datasets with a mix of continuous and categorical attributes is encountered routinely by data analysts. This work presents a method for clustering such datasets using Homogeneity Analysis. An optimal Euclidean representation of mixed datasets is obtained using Homogeneity Analysis. This representation is then clustered. The clustering solutions from this method are compared to the clustering solutions obtained using the method based on the Gower distance that is popularly used with such datasets. This comparison is made on datasets that have been the subject of other research investigations. The Homogeneity Analysis solution is an eigenvalue based solution. The eigenvalues are used to produce the optimal Euclidean representation. Even with a single eigenvalue, the Homogeneity Analysis based solution performed better than the method based on the Gower distance. Extending the solution to use multiple eigenvalues from the Homogeneity Analysis solution is illustrated on real world datasets. This method can be used in conjunction with the mini-batch K-Means algorithm to cluster large datasets. This is illustrated on a real world dataset. The relevant theory from Homogeneity Analysis is presented.

1 Introduction

Mixed datasets (data sets with a mixture of continuous and categorical attributes) are encountered routinely by data analysts. However while there is a considerable body of work on clustering techniques for continuous datasets, clustering of categorical and mixed datasets continues to be an area of active

*Chennai Mathematical Institute

research. Techniques for clustering categorical data have been studied by other researchers and solutions to this problem exist Anderlucci and Hennig [2014]. This work presents an approach to clustering mixed datasets that is based on Homogeneity Analysis Michailidis and de Leeuw [1998]. Homogeneity Analysis determines an optimal scaling of the levels of the categorical variables in a latent Euclidean space. This work shows that clustering solutions based on such a representation can be closer to the ground truth than methods commonly used to cluster mixed datasets today. This method can be applied effectively to very large datasets. An illustration of this method to several real world datasets including a large dataset (about four hundred and twenty thousand records) is provided. The rest of this paper is organized as follows. In section 2 we motivate Homogeneity Analysis and provide the relevant ideas used in developing a clustering solution. In section 3 we illustrate the difference between a method that is popularly used to cluster mixed datasets today (based on the gower distance) and the Homogeneity Analysis based solution. In section 4 we illustrate the application of Homogeneity Analysis to real world datasets. In section 5 we illustrate the application of Homogeneity Analysis to a large real world mixed dataset. In section 6 we provide the details of key `R` and `python` packages used in this work. Section 7 presents the conclusions from this work.

2 The Homogeniety Analysis Based Solution

Homogeneity Analysis is a multivariate analysis technique that determines an optimal scaling of categorical variables. The key ideas are described in this section and the reader is referred to Michailidis and de Leeuw [1998] for a detailed description of the technique.
To motivate the idea of Homogeneity Analysis, consider a dataset with categorical attributes. Such datasets do not have a natural representation in Euclidean space. We seek a procedure to represent the categorical variables of our dataset in a p dimensional Euclidean space, $\mathbb{R}^p$, in an optimal manner. Since it is only the categorical attributes for which this needed, the analysis that follows is for the subset of the dataset that is categorical. The continuous attributes are excluded. $\mathbb{D}$ represents the subset of our dataset with n rows that contains only the categorical attributes. The Homogeneity Analysis representation of the dataset is characterized by the following key elements:

1. The true representation of the row or instance in the latent Euclidean space, $\mathbb{R}^p$.

2. The optimally scaled representation of the row in terms of the observed attributes. This representation uses an optimal real value for each categorical attribute in $\mathbb{R}^p$. This is what we seek to learn.

3. An edge between an object's true representation and its approximation. This edge represents the loss of information due to the categorical nature of the object's attributes.

Such a representation induces a bipartite graph. The disjoint vertex sets for this graph are the object's true representation and its approximate attribute representation in the latent Euclidean space. This idea is represented in Figure 1

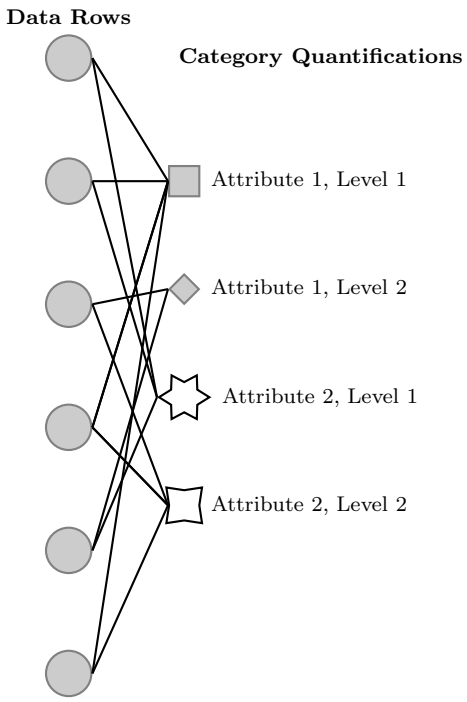

Figure 1: Representation of a Dataset with Categorical Attributes as a Bipartite Graph

The total edge length (sum of each edge length) associated with the bipartite graph corresponds to the total loss of information. An optimal representation is one that minimizes this edge length. This task can be expressed as

an optimization problem.

The following definitions provide the components of the Homogeneity Analysis solution:

Definition 1. The true representation of an element of $\mathbb{D}$ in $\mathbb{R}^p$ is called the *object score* of the element. The object scores of $\mathbb{D}$ are represented by a matrix $\mathbf{X}$ of dimension $n \times p$.

Definition 2. A categorical attribute's representation in $\mathbb{R}^p$ is called its *category quantification*. The category quantification for the attributes of $\mathbb{D}$ is represented by a matrix $\mathbf{Y}$ called the category quantification matrix. If $\mathbf{J}$ represents the attributes of $\mathbb{D}$, then the number of category quantifications, ncc, is given by:

$$ncc = \sum_{j \in \mathbf{J}} l_j$$

l_j represents the number of levels for attribute j.

The dimension of $\mathbf{Y}$ is $ncc \times p$.

Definition 3. The optimally scaled representation of $\mathbb{D}$ in $\mathbb{R}^p$ requires the use of an *indicator matrix*. The indicator matrix, $\mathbf{G}$, is a representation of $\mathbb{D}$ using a one hot encoding scheme. In a one hot encoding scheme, each attribute of $\mathbb{D}$ is represented by a set of columns corresponding to the number of levels of the attribute. The attribute level taken by the attribute for a particular row is encoded as a 1, other levels are encoded as 0. The dimension of $\mathbf{G}$ is $n \times ncc$.

Definition 4. The *optimally scaled representation* of $\mathbb{D}$ in $\mathbb{R}^p$ is expressed as product of the indicator matrix $\mathbf{G}$ and the category quantification matrix $\mathbf{Y}$.

$$\mathbb{D} = \mathbf{G}.\mathbf{Y}$$

The dimension of the optimally scaled representation is $(n \times ncc).(ncc \times p) = (n \times p)$.

Definition 5. Let J be the total number of categorical attributes. The *loss* associated with the Homogeneity Analysis based solution is the difference between the true representation and the optimally scaled representation.

The loss function can be expressed in terms of the attributes as follows:

$$\sigma = \frac{1}{J}\sum_{j=1}^{j=J} SSQ(X - G_j.Y_j) \tag{1}$$

$$= \frac{1}{J}\sum_{j=1}^{j=J} \mathbf{tr}(X - G_j.Y_j)'(X - G_j.Y_j) \tag{2}$$

Here, $SSQ(H)$ refers to sum of the squares of elements of matrix H

Homogeneity Analysis solves the following optimization problem:

$$\begin{aligned} \underset{X,Y}{\text{minimize}} \quad & \frac{1}{J}\sum_{j=1}^{j=J} \mathbf{tr}(X - G_j.Y_j)'(X - G_j.Y_j) \\ \text{subject to} \quad & \mathbf{X}^T.\mathbf{X} = n.\mathbf{I}_p \\ & u^T\mathbf{X} = 0 \end{aligned} \tag{3}$$

Where:

- $\mathbf{I}_p$ is the identity matrix of size p
- u is a vector of ones (of length n)

The above constraints standardize $\mathbf{X}$ and force the solution to be centered around the origin. The constraints also eliminate the trivial solution: $\mathbf{X} = 0$ and $\mathbf{Y} = 0$. This optimization problem is solved using an Alternating Least Squares (ALS) algorithm. A brief sketch of the steps of the algorithm is

provided in Algorithm 1

Data: **G**
Result: **X** and **Y**
1 $\mathsf{X} \leftarrow$ `InitializeX()`;
2 $D_j \leftarrow G_j'.G_j$;
4 **while** *solution has not converged* **do**
6 `/* Minimize` Y_j `based on current values of` X`. Essentially we are solving` $X = G_j.Y_j + \epsilon$ `for all` $j \in \mathbf{J}$`. The solution for this is given below */`
7 $Y_j \leftarrow D_j^{-1} G_j' X$;
9 `/* Now fix` Y_j `and minimize` X`. The optimal value of` X `is given below */`
10 $\mathsf{X} \leftarrow J^{-1} \sum_{j=1}^{j=J} G_j Y_j$;
12 `/* Center and Orthonormalize X so that the constraints are satisfied. Orthornormalization is performed by an algorithm like Gram Schmidt */`
13 $\mathsf{X} \leftarrow$ `CenterAndOrthonormalizeX()`;
14 **end**

Algorithm 1: Summary of the ALS Algorithm for Homogeneity Analysis

The Homogeneity Analysis problem can be expressed as an eigenvalue problem (see Michailidis and de Leeuw [1998] for details). The loss function of the Homogeneity Analysis solution (Equation 2) can be expressed in terms of the eigenvalues of the average projection matrix $\mathbf{P}^*$, for the subspace spanned by the columns of the indicator matrix $\mathbf{G}$ (see Michailidis and de Leeuw [1998]). For attribute j, the projection matrix is $P_j = G_j.D_j^{-1}.G_j$. The average projection matrix for J attributes is given by $\mathbf{P}^* = J^{-1} \sum_{j=1}^{j=J} P_j$. The loss σ, for the Homogeneity Analysis solution can be expressed in terms of the eigenvalues of $\mathbf{P}^*$. (see Michailidis and de Leeuw [1998]):

$$\sigma = n \left(p - \sum_{s=1}^{s=p} \lambda_s \right) \tag{4}$$

λ_s represents the eigenvalues of $\mathbf{P}^*$

An inspection of Equation 4 shows that the number of eigenvalues to use with Homogeneity Analysis solution (p in Equation 4) is a parameter to be chosen. Increasing the number of eigenvalues decreases the loss, however this also increases the dimensionality of the category quantification (each categorical attribute level has a p dimensional representation). In this study we found

that using the first eigenvalue ($p = 1$) alone to determine the optimal real valued representation yielded good clustering solutions in many cases. This is consistent with the fact that the first eigenvalue holds most information about the attributes real valued representation. If we use higher values of p (i.e, $p > 1$), then we would replace each categorical value in our original dataset by a p tuple of values. The Homogeneity Analysis solution produces a scree plot of the eigenvalues which can be used to determine the number of eigenvalues to use for a particular problem. This is illustrated in the sections illustrating the application of Homogeneity Analysis (see section 4).
Van Buuren and Heiser [1989] had used Homogeneity Analysis to determine a clustering solution. This solution goes by the name of GROUPALS. The approach used in this work is different from that used by GROUPALS. GROUPALS solves the clustering problem (partitioning n objects into k groups) and the problem of determining an optimal scaling concurrently. GROUPALS uses minimizing the sum of the square errors as the objective to determine a clustering. In this work, the clustering problem is decoupled from the optimal scaling problem. This decoupling permits us to choose any criterion (for example, the average silhouette width) to determine an optimal clustering. The GROUPALS approach of Van Buuren and Heiser [1989] also requires the number of clusters to be specified as an input to the algorithm. This may not necessarily be known at the time of performing the analysis. When the problems are decoupled, the optimal scaling solution can be used with any one of a variety of clustering quality metrics like silhouette width or within group sum of the squares etc., to determine an optimal number of clusters. The final clustering solution can be obtained by computing a clustering solution on the optimally scaled data with the optimum number of clusters.

3 Homogeneity Analysis Versus Gower Distance Based Clustering

To illustrate the utility of a Homogeneity Analysis based solution we compare a ground truth clustering solution with the following:

1. The clustering solution obtained from a popular method used to cluster mixed datasets.

2. The clustering solution obtained using Homogeneity Analysis

This comparison is performed on several datasets. The following are needed to make this comparison:

1. The ground truth

2. A metric to compare clustering solutions.

Partitioning Around Mediods (PAM) using the Gower distance (see Gower [1971]) is a very common method used to cluster mixed datasets. We compare the clustering solution produced by this method to the clustering solutions produced by Partitioning Around Mediods using the dataset obtained from Homogeneity Analysis.
Accordingly, the material in this section is organized as follows. In section 3.1 we present the metric used to compare clustering solutions. In section 3.2 we present the methodology used to establish the ground truth for the datasets used in this study.

3.1 Metric to Compare Clustering Solutions

The metric used to compare clustering solutions is the Variation of Information distance (see Meilă [2007]). This metric is based on information theory and requires the following definitions:

Definition 6. A clustering $\mathcal{C}$ produces a set of k clusters, $\{C_1, C_2, \cdots, C_k\}$. The *probability* that an element of the dataset is assigned to a cluster C_i is given by:

$$P(i) = \frac{|C_i|}{n} \quad i \ = \ \{1, 2, \cdots, k\}$$

$|C_i|$ represents the number of elements belonging to cluster C_i.

Definition 7. The *entropy* associated with a clustering $\mathcal{C}$ is then given by:

$$\mathcal{H}(\mathcal{C}) = \sum_{i=1}^{i=k} P(i) \log_2 P(i)$$

Wagner (see Wagner and Wagner [2007]) describes entropy as an informal measure of the uncertainty associated with a cluster assignment of a random element of the dataset. A natural extension of the idea of entropy is that of mutual information.

Definition 8. Given two clusterings, $\mathcal{C}$ and $\mathcal{C}'$, *mutual information* provides a measure of the reduction of uncertainty associated with the cluster assignment of a random element in clustering $\mathcal{C}$ given that we know its cluster

assignment in $\mathcal{C}'$.
The mutual information for a given pair of clusterings $\mathcal{C}$ and $\mathcal{C}'$ is given by:

$$\mathcal{I}(\mathcal{C},\mathcal{C}') = \sum_{i=1}^{i=k}\sum_{j=1}^{j=1} P(i,j).\log_2 \frac{P(i,j)}{P(i).P(j)}$$

$P(i,j)$ is the probability that a data instance belongs to cluster C_i in clustering $\mathcal{C}$ and to cluster C_j in clustering $\mathcal{C}'$. It is calculated as:

$$P(i,j) = \frac{|C_i \cap C_j|}{n}$$

$|C_i \cap C_j|$ represents the number of elements that are in both C_i and C_j.

We are now ready to define the Variation of Information distance to compare a pair of clustering solutions.

Definition 9. The *Variation of Information* (see Meilă [2007])distance between a pair of clustering solutions $\mathcal{C}$ and $\mathcal{C}'$ is given by:

$$\mathcal{VI}(\mathcal{C},\mathcal{C}') = \big(\mathcal{H}(\mathcal{C}) - \mathcal{I}(\mathcal{C},\mathcal{C}')\big) + \big(\mathcal{H}(\mathcal{C}') - \mathcal{I}(\mathcal{C},\mathcal{C}')\big) \tag{5}$$

Wagner (see Wagner and Wagner [2007]) describe the first term of Equation 5 as a measure of the amount of information about $\mathcal{C}$ that we loose while we go from clustering $\mathcal{C}$ to clustering $\mathcal{C}'$. The second term of Equation 5 represents the amount of information that we still have to gain about $\mathcal{C}'$, as we go from clustering $\mathcal{C}$ to clustering $\mathcal{C}'$.

3.2 Establishing the Ground Truth

We establish the ground truth solution for the data sets used in this study by careful construction. The datasets used to establish the ground truth are all numerical datasets that have been studied by other researchers. The number of clusters in these datasets is known. We discretize a subset of these attributes and use the resulting dataset for our study. The discretization is achieved by replacing each value of the attribute being discretized by the quartile it belongs to. The categorical variables obtained by discretization are all ordinal. The clustering solutions obtained using the dataset prior to discretization are designated as the ground truth. The clustering solutions using the PAM algorithm on the non-discretized dataset is designated as the ground truth. We designate this solution as the ground truth because this is the solution we would have come up with if we had access to the dataset

where all the attributes are continuous.

The procedure to analyze the datasets is summarized in Algorithm 2.

```
Data: A dataset with continuous attributes
Result: The VI distance between clustering solutions for discretized
        dataset and the ground truth
/* the clustering algorithms below are run on the dataset with
   numeric attributes                                          */
ground.truth.distance.based ← Cluster.Data.using.PAM();
/* a few attributes of the numeric dataset are discretized by
   representing them by their quantiles                        */
discretized.dataset ← Discretize.Some.Attributes();
/* The clustering algorithm below is run on the discretized
   dataset using a dissimilarity matrix computed based on the
   Gower distance                                              */
discretized.PAM.clustering.solution. ← Cluster.Data.using.PAM();
/* Compute the homals based dataset                            */
homals.based.dataset ← Compute.Homals.Based.Dataset();
/* The algorithms below are run on the homals based dataset    */
homals.PAM.clustering.solution ← Cluster.Data.using.PAM();
/* Compare the clustering solutions obtained from Gower
   distance based clustering and Homogeneity Analysis based
   clustering to the ground truth                              */
Compare.Clustering.Solutions.With.Ground.Truth();
```

Algorithm 2: Procedure for Real Datasets

The reader might wonder why the category quantifications are used instead of the object scores in comparing clustering solutions. For this part of the study we want to compare the ground truth clustering solution to the clustering solutions obtained using the Gower distance and the dataset we obtain from Homogeneity Analysis. To make this comparison the datasets for each of the cases being compared must have the same attributes. Since our objective is to make a comparison to other clustering solutions we choose the category quantifications instead of the object scores in determining a numerical representation for our dataset. Using the category quantifications also permits us to profile the clustering solution. We can examine the ranges of values of the attributes in each cluster to obtain insights from it. Category quantifications were used to for the numerical representation of the categorical variables in this study.

3.3 Datasets

The datasets used for this part of the study have been studied by other researchers. The number of clusters for these datasets are known. The following datasets were used for this study:

1. The Iris dataset is attributed to Edgar Anderson Anderson [1935] and has been studied extensively. It is commonly accepted that this dataset has three clusters.

2. The Wine dataset was obtained from the University of California Machine Learning Repository (UCIMLR) Lichman [2013]. The dataset captures the results of chemical analysis performed on wines from a particular region in Italy from three different growers. Bilenko et al. [2004] report three clusters for this dataset.

3. The seeds dataset was obtained from the UCIMLR. This dataset captures the geometric characteristics of three different types of wheat kernels. It has been studied by Charytanowicz et al. [2010] who report three clusters.

4. The glass dataset was obtained from the UCIMLR. This dataset captures the properties of different types of glass obtained from forensic investigations. Eick et al. [2004] report six clusters for this dataset.

5. The swiss Banknote dataset was obtained from the *mclust* R package Fraley et al. [2012]. Atkinson and Riani [2007] report three clusters for this dataset.

3.4 Results - Homogeneity Analysis Versus Gower Distance Based Clustering

The procedures described above compare a ground truth clustering solution with clustering solutions produced by methods used today and the clustering solutions produced using Homogeneity Analysis. Clustering solutions are compared in terms of Variation of Information distance. The results are shown in Table 1 An examination of the results in Table 1 shows that the clustering solution obtained using Homogeneity Analysis is closer to the ground truth clustering solution than the one obtained using the Gower distance for all datasets. All these solutions used only the first eigenvalue of the Homogeneity Analysis solution. Even with such a representation, the clustering solutions using Homogeneity Analysis are closer to the ground

Dataset	PAM	Homals PAM
Iris	1.029	**0.797**
Wine	0.924	**0.481**
Banknote	1.174	**0.948**
Seeds	0.579	**0**
Glass	2.919	**1.433**

Table 1: Real Datasets, VI Distance from Ground Truth

truth. Examples of using more than one eigenvalue are illustrated in the section 4.

4 Application of Homogeniety Analysis to Clustering Real World Datasets

In this section we illustrate the application of Homogeneity Analysis to real world datasets. The use of multiple eigenvalues from the Homogeneity Analysis solution is also illustrated.

4.1 The mtcars Dataset

The mtcars dataset is a small dataset that contains mixed data. It is a small dataset that is available in the R datasets package R Core Team [2015]. This dataset was picked because it contains mixed data and is small enough to illustrate the clusters well. The dataset was extracted from the 1974 US Motor Trend magazine. It contains fuel consumption information as well as design features of a small (32) set of cars. Table 2

Attribute	Description	Type
mpg	Miles Per Gallon	Continuous
cyl	Number of cylinders	Nominal
disp	Displacement	Continuous
hp	Horse Power	Continuous
drat	Real Wheel Axle Ratio	Continuous
Wt	Weight	Continuous
qsec	Time for 0.25 mile	Continuous
vs	V/S	Nominal
am	Transmission type	Nominal
gear	Number of forward gears	Ordinal
carb	Number of carburetors	Ordinal

Table 2: mtcars Dataset Description

Homogeneity Analysis was used to obtain a numeric representation for the categorical variables in this dataset. The resulting dataset was then clus-

tered. The number of clusters were determined using the average silhouette width criterion Rousseeuw [1987]. The number of eigenvalues to use for the numerical representation is obtained from the scree plot associated with Homogeneity Analysis solution (see Figure 2). An examination of Figure 2 shows that four eigenvalues are appropriate for this dataset. The clustering results are shown in Table3.

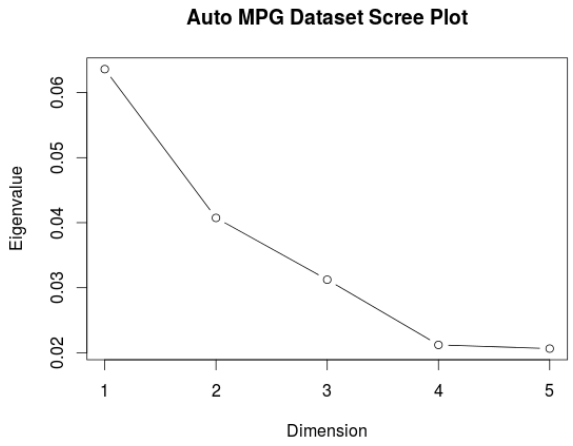

Figure 2: Scree Plot for the Homogeneity Analysis Solution for mtcars

Cluster	Car	HP	Cylinders
1	Mazda RX4	110.00	6
1	Mazda RX4 Wag	110.00	6
1	Datsun 710	93.00	4
1	Merc 240D	62.00	4
1	Merc 230	95.00	4
1	Merc 280	123.00	6
1	Merc 280C	123.00	6
1	Fiat 128	66.00	4
1	Honda Civic	52.00	4
1	Toyota Corolla	65.00	4
1	Toyota Corona	97.00	4
1	Fiat X1-9	66.00	4
1	Porsche 914-2	91.00	4
1	Lotus Europa	113.00	4
1	Volvo 142E	109.00	4

Cluster	Car	HP	Cylinders
2	Hornet 4 Drive	110.00	6
2	Hornet Sportabout	175.00	8
2	Valiant	105.00	6
2	Duster 360	245.00	8
2	Merc 450SE	180.00	8
2	Merc 450SL	180.00	8
2	Merc 450SLC	180.00	8
2	Cadillac Fleetwood	205.00	8
2	Lincoln Continental	215.00	8
2	Chrysler Imperial	230.00	8
2	Dodge Challenger	150.00	8
2	AMC Javelin	150.00	8
2	Camaro Z28	245.00	8
2	Pontiac Firebird	175.00	8
2	Ford Pantera L	264.00	8
2	Ferrari Dino	175.00	6
2	Maserati Bora	335.00	8

Table 3: Cluster Profiles for the mtcars Dataset

An examination of Table 3 shows that cars with similar characteristics cluster together. Cluster 1 captures smaller cars with mostly four and some six cylinders. Cluster 2 captures more powerful six and eight cylinder cars

4.2 The statlog (Heart) Dataset

The Heart dataset from the UCI ML repository Lichman [2013] was used for this study. It is a dataset with mixed attributes of people with and without heart disease. Homogeneity Analysis was used to determine the numerical representation of the categorical variables in this dataset. A description of the dataset is shown in Table4

Variable	Description	Type
Age	Age	Continuous
Sex	Sex	Nominal
CP	Chest Pain Type	Nominal
RBP	Resting Blood Pressure	Continuous
SC	Serum Cholesterol	Continuous
FBS	Is fasting blood Sugar over 120	Nominal
RECG	Resting electrocardiographic results (values 0,1,2)	Nominal
MHR	Maximum heart rate achieved	Continuous
EIA	Exercise induced angina	Nominal
OP	Oldpeak = ST depression induced by exercise relative to rest	Continuous
SST	The slope of the peak exercise ST segment	Ordinal
NV	Number of major vessels (0-3) colored by flourosopy	Continuous
Thal	Thal: 3 = normal; 6 = fixed defect; 7 = reversable defect	Nominal

Table 4: Heart Disease Data Types

The number of eigenvalues to use for the numerical representation is obtained from the scree plot obtained from the Homogeneity Analysis solution (see Figure 3). The scree plot shows that using six eigenvalues will capture all the information required for a good numerical representation of the categorical attributes. The resulting numerical dataset was clustered. The number of clusters were determined using the average silhouette width criterion Rousseeuw [1987].
A summary view of the cluster profiles is shown in Table 5. Two clusters were most appropriate for this dataset based on the average silhouette width criterion. Cluster 2 is a relatively disease free cluster with 85 percent of the subjects having no disease. Cluster 1 has a much higher proportion of subjects with disease (about 61 percent - $\frac{105}{170}$).

Cluster	Count	Disease Status
1	65	Absent
1	105	Present
2	85	Absent
2	15	Present

Table 5: Heart Disease Clusters

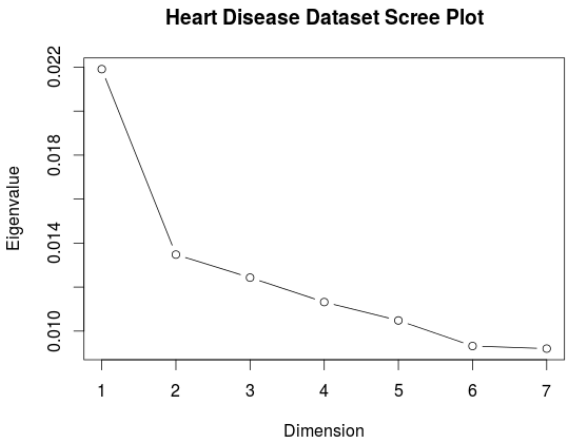

Figure 3: Scree Plot for the Homogeneity Analysis Solution for Heart Disease

5 Application to Big Data

In this section the details of applying Homogeneity Analysis to clustering a big data set is provided. The details of the dataset used for the study are provided in section 5.1. The methodology used for clustering mixed big datasets is described in section 5.2. The results from the experiments are described in section 5.3.

5.1 The Dataset

The US Department of Transportation publishes monthly reports of airline on time performance USDOT [2016]. The data used for this study represents the on time arrival performance for the January 2016 time period. The dataset has about four hundred and forty five thousand records including some outliers. Records within the 99 th percentile of arrival delays (155 minutes) were used for this study. This (large) subset was about four hundred and twenty seven thousand records. The attributes and their description for this dataset is provided in Table 6

5.2 The Methodology

The methodology used to apply Homogeneity Analysis to cluster big datasets with a mixture of continuous and categorical variables is based on using a random sample obtained using stratified sampling. For this dataset the ori-

	Attribute	Type	Description
1	DAY_OF_MONTH	Ordinal	Day of January for flight record
2	DAY_OF_WEEK	Ordinal	Day of week for flight record
3	CARRIER	Nominal	Carrier (Airline) for the flight record
4	ORIGIN	Nominal	Origin airport code
5	DEST	Nominal	Destination airport code
6	DEP_DELAY	Continuous	Departure Delay in minutes
7	TAXI_OUT	Continuous	Taxi out time in minutes
8	TAXI_IN	Continuous	Taxi in time in minutes
9	ARR_DELAY	Continuous	Arrival delay in minutes
10	CRS_ELAPSED_TIME	Continuous	Flight duration
11	NDDT	Continuous	Departure time in minutes from midnight January 1 2016

Table 6: Delay Data January 2016

gin and destination variables were picked as the stratification variables. This ensures all origin destination pairs are represented in the sample. Homogeneity Analysis is performed on this sample and the number of eigenvalues to be used is based on the analysis of the scree plot obtained from the solution. The category quantifications obtained from the sample are applied to the big dataset. This is a simple recoding procedure that involves replacing each level of a categorical attribute with the set of values corresponding to that level obtained from the Homogeneity Analysis solution. The number of values with which each level of the categorical variable is replaced with depends on the number of eigenvalues we pick for the solution. Since the dataset is large, the mini-batch K-Means algorithm is used to cluster the dataset. Mini-batch K-Means is a variant of the regular K-Means algorithm (see Sculley [2010]). The algorithm consists of two steps. In the first step a batch of samples is picked from the large dataset and is clustered using the regular K-Means algorithm. In the next step another batch of samples are picked. Each sample is assigned to the nearest centroid obtained from the previous step and the cluster centroids are updated. The elbow plot method is used to pick the number of clusters for the clustering solution. Mini-batch K-Means is run for a range of values for the number of clusters(K). The mini-batch K-Means solution calculates a quantity called Inertia that represents the total sum of the square distances between the points in the dataset to their respective centroids. The optimal number of clusters is one where there is no significant change in the Inertia with an increase in the number of clusters. This plot has a characteristic elbow shape and the point where there is no significant change in Inertia with an increase in the number of clusters usually appears at the elbow (see section 5.3).
Once the optimal number of clusters has been determined using the elbow plot, the clustering solution for analysis is obtained by running mini-batch K-Means with the optimal number of clusters on the dataset obtained from the recoding procedure. For experiments in this study a batch size of 2000

was used. Principal Component Analysis(PCA) is used to visualize the clustering solution. Incremental Principal Component Analysis (IPCA) is used instead of PCA because of the large dataset size (Ross et al. [2008]). The method is based on building a low rank approximation to the dataset.

5.3 Discussion of Results

As discussed in section 5.2, the first step in the solution is to pick a sample for Homogeneity Analysis. For this dataset, the $ORIGIN$ and $DESTINATION$ (see Table 6) attributes were the stratification variables. The number of records for each origin and destination pair is proportional to the number of records for the same pair in the population. A two percent sample was picked. This amounts to a sample size of about ten thousand records, which the homogeneity analysis implementation (see de Leeuw and Mair [2009]) could easily handle.
The scree plot produced by the Homogeneity Analysis solution is shown in Figure 4 Three eigenvalues were used for the solution. The level of each

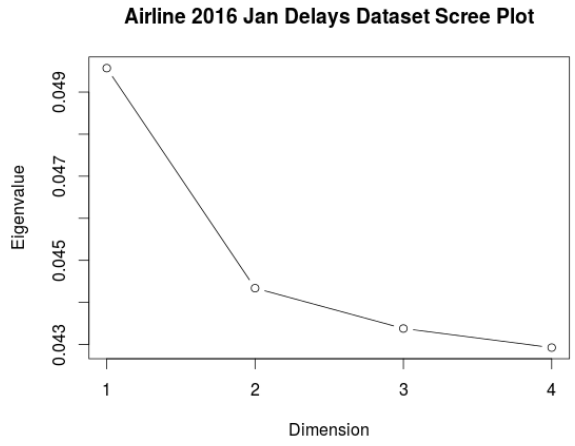

Figure 4: *Scree Plot for the January 2016 Delay Data*

categorical variable were replaced with the category quantifications for that level obtained from the Homogeneity Analysis solution. Since three eigenvalues were picked, each level is replaced with a set of three numeric values. After this recoding is completed, the dataset is now ready for clustering analysis. The next task is to determine the number of clusters that are optimal for this dataset. This is obtained using an elbow plot that plots the

Inertia versus the number of clusters (see section 5.2). This plot is shown in Figure 5. An analysis of Figure 5 shows that 35 clusters are appropriate for

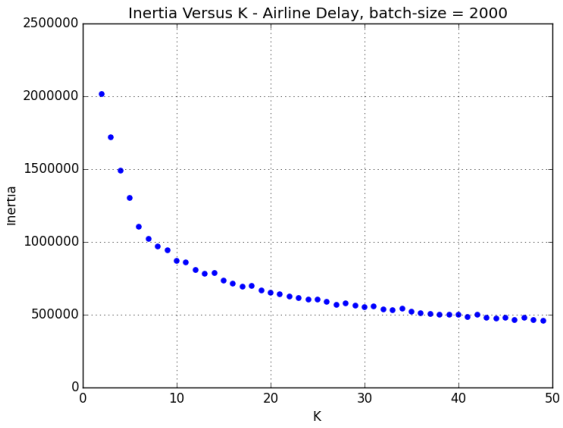

Figure 5: *Elbow Plot for the January 2016 Delay Data*

this dataset. The mini-batch K-Means solution can now be computed with 35 clusters. An Incremental Principal Component Analysis implementation (Pedregosa et al. [2011]) was used to visualize the computed solution on the first two principal components. This is shown in Figure 6 The clustering

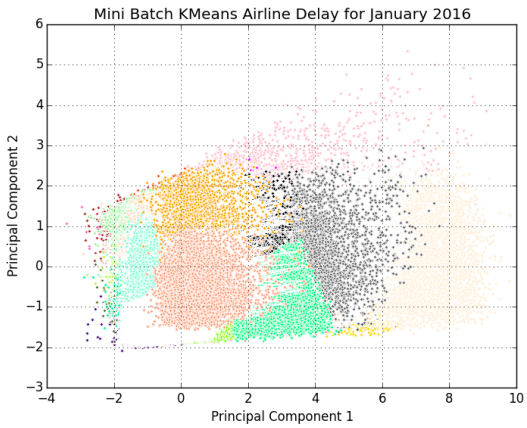

Figure 6: *Cluster Visualization for the January 2016 Delay Data*

solution is now ready to be profiled and analyzed. The salient facts from such an analysis are presented below. The distribution of cluster sizes is

shown in Figure 7 An analysis of Figure 7 shows that we have many small

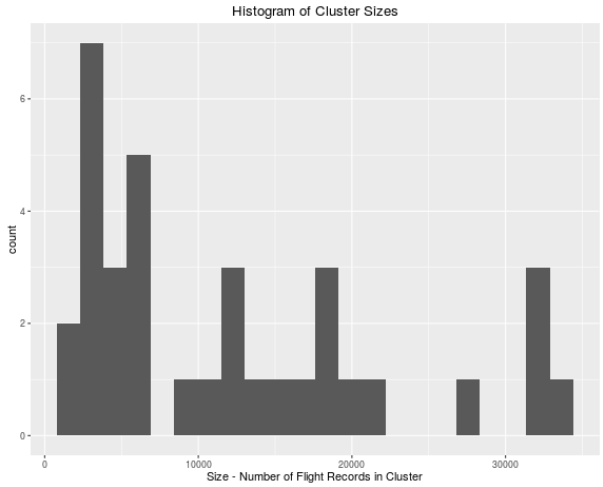

Figure 7: *Histogram of Cluster Size (Number of Flight Records)*

sized clusters (less than ten thousand flight records), some medium sized clusters (between eight thousand and twenty two thousand flight records) and a few large clusters (greater than twenty five thousand flight records).

The total arrival delay for a cluster and the average arrival delay for a cluster are of interest to us. Histograms of these (with a density plot overlay) are shown in Figure 8 and Figure 9. The clusters to the extremes of

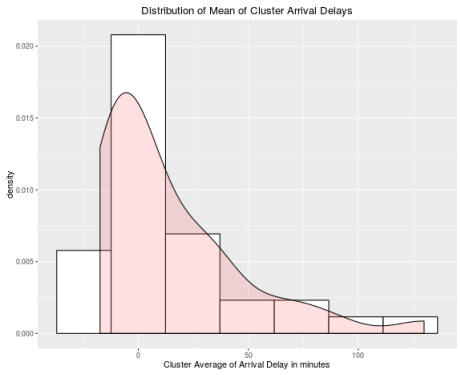

Figure 8: *Histogram Cluster Mean of Delay*

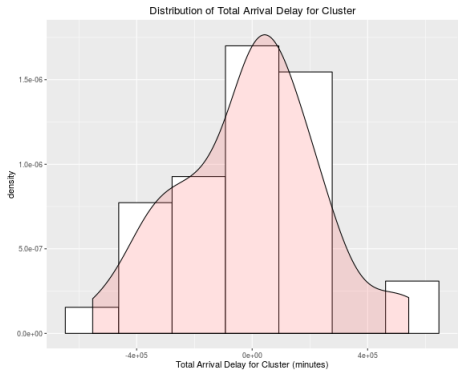

Figure 9: *Histogram - Total Delay for Cluster*

Figure 8 and Figure 9 are of interest to us. The extreme right - arrival delays are of more interest. A brief profile of the cluster to the extreme

right of Figure 8 is as follows. The cluster to the extreme right of Figure 8 is the cluster with the highest average arrival delay. The records in the cluster were aggregated by the origin of the flight and the total arrival delay for that origin was computed. The results for the origin airports with the highest totals (top five) of arrival delays are shown in Table 7

Origin Airport Code	Airport	Total Arrival Delay (min)	Number of Records
ORD	O'Hare Intl. Airport (Chicago)	31532	244
SFO	San Franscisco Intl. Airport	28723	221
LAX	Los Angeles Intl. Airport	21621	166
ATL	Hartsfield Intl Airport (Atlanta)	20678	156
LAS	McCarran Intl Airport (Las Vegas)	19513	145

Table 7: Top Five Origin Airports for Delayed Cluster

Similarly we can aggregate the data for the delayed cluster by the day of week and compute the total arrival delay for each day of the week. The results are shown in Table 8

Day of Week	Total Arrival Delay (min)	Number of Records
7	90133	694
5	86838	664
1	67054	522
2	61476	475
3	57561	440
4	55767	428
6	54601	422

Table 8: Total Delays by Day of Week for Delayed Cluster

Table 7 shows that the worst delays for the delayed cluster are associated with the major airports. Table 8 shows that the worst delays for the delayed cluster are associated with Sunday, Friday and Monday.

6 Implementation Details

The experiments conducted for this study used R and python for their implementation. R was used for Homogeneity Analysis and for most of the experiments. Python was used for the big data experiments conducted in this study. The details of the tools are given below.

6.1 R Tools

The following are the key R packages and methods used for this study:

1. The `pamk` method from the `fpc` package was used to determine the distance based clustering solutions for this study. This method supports determining the number of clusters using a variety of criteria. The

average silhouette width criterion was used to determine the number of clusters for this study.

2. The `daisy` method from the `cluster` package was used to compute the dissimilarity matrix for mixed datasets. Clustering the dissimilarity matrix using a method like `pamk` is a popular method of using a distance based clustering method on mixed datasets. The gower distance is used as the distance metric for the `daisy` method.

3. The `homals` method from the `homals` package was used to perform Homogeneity Analysis on the datasets used in this study.

4. The `vi.dist` method from the `mcclust` package was used to compute the Variation of Information distance between two clustering solutions.

5. The `strata` method from the `sampling` package was used to obtain a stratified sample for the big data experiments reported in this study.

6.2 Python Tools

The `scikit-learn`Pedregosa et al. [2011] python package was used for the big data experiments reported in this study. The details of the methods used are as follows:

1. `MiniBatchKMeans` method was used for the mini-batch K-Means experiments and for the final clustering solution.

2. `IncrementalPCA` method was used for the visualization of the clustering solution obtained using mini-batch K-Means.

7 Conlusion

Datasets with categorical attributes are commonly encountered by data analysts. Homogeneity Analysis provides a theoretical basis to determine a Euclidean representation of such attributes. Such a representation immediately puts a wide array of algorithms at the disposal of the analyst. This study focuses on clustering such datasets. Clustering a dissimilarity matrix based on the Gower distance is a common way to cluster mixed datasets. Homogeneity Analysis has an eigenvalue based solution. Experiments on datasets that have been studied by other researchers showed that even with a single eigenvalue, the Homogeneity Analysis based solution performs better than the method commonly used today. Since the dissimilarity matrix

calculation involves the calculation of the pairwise dissimilarities of all data instances, this method is applicable only to small or moderate sized datasets. As was illustrated in this study, sampling and Homogeneity Analysis can be used in conjunction with the mini-batch K-Means algorithm to cluster large mixed datasets.
GROUPALS Van Buuren and Heiser [1989] solves a similar problem. However it solves the problem of determining an optimal numerical representation and the clustering solution concurrently. It also uses the squared error minimization as the loss function in its solution. It requires the analyst to know the number of clusters in the data. In this work, the clustering solution and the problem of determining an optimal numerical representation are decoupled. The Homogeneity Analysis solution is used as the basis for the clustering solution. This decoupling permits the analyst to apply any criteria (such as the average silhouette width Rousseeuw [1987]) to determine the optimal number of clusters, and then compute an optimal clustering solution.

Acknowledgements

The author would like to thank Dr. Sourish Das (faculty member at the Chennai Mathematical Institute) and Dr. Sasanka Roy (faculty member at the Indian Statistical Institute, Kolkata) for reviewing this manuscript and providing many insightful suggestions and comments.

References

Anderlucci, L. and Hennig, C. (2014). The clustering of categorical data: A comparison of a model-based and a distance-based approach. *Communications in Statistics - Theory and Methods*, 43(4):704–721.

Anderson, E. (1935). The Irises of the Gaspe Peninsula. *Bulletin of the American Iris Society*, 59:2–5.

Atkinson, A. C. and Riani, M. (2007). Exploratory tools for clustering multivariate data. *Computational Statistics & Data Analysis*, 52(1):272–285.

Bilenko, M. et al. (2004). Integrating constraints and metric learning in semi-supervised clustering. In *Proceedings of the twenty-first international conference on Machine learning*, page 11. ACM.

Charytanowicz et al. (2010). Complete gradient clustering algorithm for features analysis of x-ray images. In *Information technologies in biomedicine*, pages 15–24. Springer.

de Leeuw, J. and Mair, P. (2009). Gifi methods for optimal scaling in R: The package homals. *Journal of Statistical Software*, 31(4):1–20.

Eick, C. F. et al. (2004). Using representative-based clustering for nearest neighbor dataset editing. In *Data Mining, 2004. ICDM'04. Fourth IEEE International Conference on*, pages 375–378. IEEE.

Fraley, C. et al. (2012). *mclust Version 4 for R: Normal Mixture Modeling for Model-Based Clustering, Classification, and Density Estimation.*

Fraley, C. and Raftery, A. E. (2002). Model-based clustering, discriminant analysis and density estimation. *Journal of the American Statistical Association*, 97:611–631.

Fritsch, A. (2012). *mcclust: Process an MCMC Sample of Clusterings.* R package version 1.0.

Gower, J. C. (1971). A general coefficient of similarity and some of its properties. *Biometrics.*

Hennig, C. (2015). *fpc: Flexible Procedures for Clustering.* R package version 2.1-10.

Hubert, L. and Arabie, P. (1985). Comparing partitions. *Journal of classification*, 2(1):193–218.

Jain, A. and Dubes, R. (1988). *Algorithms for clustering data.* Prentice-Hall Advanced Reference Series. Prentice Hall PTR.

Kaufman, L. and Rousseeuw, P. (2005). *Finding Groups in Data: An Introduction to Cluster Analysis.* Wiley Series in Probability and Statistics. Wiley.

Kim, H. (2012). *discretization: Data preprocessing, discretization for classification.* R package version 1.0-1.

Lichman, M. (2013). UCI machine learning repository.

Maechler, M., Rousseeuw, P., Struyf, A., Hubert, M., and Hornik, K. (2015). *cluster: Cluster Analysis Basics and Extensions.*

Meilă, M. (2007). Comparing clusterings - an information based distance. *Journal of multivariate analysis*, 98(5):873–895.

Michailidis, G. and de Leeuw, J. (1998). The gifi system of descriptive multivariate analysis. *STATISTICAL SCIENCE*, 13:307–336.

Morey, L. C. and Agresti, A. (1984). The measurement of classification agreement: an adjustment to the rand statistic for chance agreement. *Educational and Psychological Measurement*, 44(1):33–37.

Pedregosa, F. et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Ramey, J. A. (2012). *clusteval: Evaluation of Clustering Algorithms*. R package version 0.1.

Ross, D. A. et al. (2008). Incremental learning for robust visual tracking. *International Journal of Computer Vision*, 77(1-3):125–141.

Rousseeuw, P. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.*, 20(1):53–65.

Sculley, D. (2010). Web-scale k-means clustering. In *Proceedings of the 19th international conference on World wide web*, pages 1177–1178. ACM.

Till, Y. and Matei, A. (2015). *sampling: Survey Sampling*. R package version 2.7.

USDOT, B. (2016). RITA airline delay data download.

Van Buuren, S. and Heiser, W. J. (1989). Clustering n objects into k groups under optimal scaling of variables. *Psychometrika*, 54(4):699–706.

Wagner, S. and Wagner, D. (2007). Comparing Clusterings – An Overview. Technical Report 2006-04, Universität Karlsruhe (TH).

Warnes, G. R., Bolker, B., and Lumley, T. (2015). *gtools: Various R Programming Tools*. R package version 3.5.0.

Wickham, H. (2011). The split-apply-combine strategy for data analysis. *Journal of Statistical Software*, 40(1):1–29.